

# An Enterprise Information Architecture: A Case Study for Decentralized Organizations

*R.W. Watson*

*This article was submitted to  
Hawaii International Conference on System Sciences  
Maui, HI  
January 4-7, 2000*

**U.S. Department of Energy**

Lawrence  
Livermore  
National  
Laboratory

**June 15, 1999**

## DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

This report has been reproduced  
directly from the best available copy.

Available to DOE and DOE contractors from the  
Office of Scientific and Technical Information  
P.O. Box 62, Oak Ridge, TN 37831  
Prices available from (423) 576-8401  
<http://apollo.osti.gov/bridge/>

Available to the public from the  
National Technical Information Service  
U.S. Department of Commerce  
5285 Port Royal Rd.,  
Springfield, VA 22161  
<http://www.ntis.gov/>

OR

Lawrence Livermore National Laboratory  
Technical Information Department's Digital Library  
<http://www.llnl.gov/tid/Library.html>

**An Enterprise Information Architecture:  
A Case Study for Decentralized Organizations**

Richard W. Watson L-66  
Chief Scientist – Computation Directorate  
Lawrence Livermore National Laboratory  
P.O. Box 808  
Livermore California, 94550  
(925) 422-9216  
[dwatson@llnl.gov](mailto:dwatson@llnl.gov)



# **An Enterprise Information Architecture: A Case Study for Decentralized Organizations**

Richard W. Watson

## **Abstract**

As enterprises become increasingly information based, making improvements in their information activities is a top priority to assure their continuing competitiveness. A key to achieving these improvements is developing an Enterprise Information Architecture (EIA). An EIA can be viewed as a structured set of multidimensional interrelated elements that support all information processes. The current ad hoc EIAs in place within many enterprises can not meet their future needs because of a lack of a coherent framework, incompatibilities, missing elements, few and poorly understood standards, uneven quality and unnecessary duplications. This paper discusses the EIA developed at Lawrence Livermore National Laboratory as a case study, for other information based enterprises, particularly those with decentralized and autonomous organization structures and cultures. While the architecture is important, the process by which it is developed and sustained over time is equally important. This paper outlines the motivation for an EIA and discusses each of the interacting elements identified. It also presents an organizational structure and processes for building a sustainable EIA activity.

## **1. Introduction**

Lawrence Livermore National Laboratory (LLNL) is an applied science laboratory whose primary missions are in national security, energy and environment, and bioscience and healthcare. Each of

these involves many sub areas of, interrelated science and technology. These activities and the business processes that support them depend centrally on the use, creation, sharing and exchange of information. It is estimated that LLNL spends approximately one quarter of its billion dollar per year budget on computer, communications, and information activities. While many of these are world-class, improvements can provide strategic advantages, increase effectiveness, and reduce costs.

A key to achieving these improvements is developing an Enterprise Information Architecture (EIA). An EIA provides the framework for planning and implementing a rich, standards-based, digital information infrastructure with well-integrated services and activities. Such an architecture, based on a shared vision and principles, can help foster:

- Easier information sharing and exchange
- Improved security and privacy
- Easier and faster building of information services tailored to particular business or programmatic needs rather than having to create part or all of the infrastructure and services from scratch
- Lower costs for system support, training, and deployment
- Increasingly effective matrix organization structure because of the common informative services, resources, and tools widely understood
- Faster and more effective response to existing and new customer requirements through these common services, resources, and tools
- Easier sharing with collaborators outside the Laboratory through wider use of industry standards
- Easier incorporation of outside vendors within chains of needed capabilities Better integration with the rest of the DOE Complex, academic community, and industry
- Lower overall institution-wide EIA-related costs.

Cook [1] points out that a prime purpose of an EIA is as a framework to develop standards and that resistance to developing and implementing an EIA and its associated standards is to be expected.

Because LLNL is a diverse decentralized organization with a long history of autonomy among its major organizational elements, the approach to successful EIA acceptance must: (1) provide an easy to understand representation of both the architecture and its requirements and (2) involve involvement from across the laboratory from start to finish. An important contribution of this work is in meeting these objectives. When examining the current *ad hoc* EIA the Chief Information Officer (CIO)

recognized that there were several challenges to overcome:

- There is no overall group or process looking at an EIA
- There are few standards
- There are no mechanisms to tap the breadth and depth of knowledge and skills on which an EIA depends
- Changes in digital technology occur faster than our ability to bring new products and services to the workplace.

To develop the base EIA an Enterprise Information Architecture Committee (IAC), with representation from all major organizations, was established by the CIO, under the author's chairmanship, with four goals:

1. Develop an EIA model based on a shared vision and principles that can serve as the framework for EIA planning and standards setting.
2. Describe (a) the current state of the Laboratory's *ad hoc* EIA, (b) the desired future EIA that meets the vision and principles, and (c) the gap between them.
3. Recommend activities that can begin closing the gap.
4. Recommend an EIA stewardship organization and process that implements the model, sustains its evolution across decades, and involves the labwide community.

## **The LLNL View of an Enterprise Information Architecture**

Developing an EIA is widely recognized as an important activity and numerous papers and other material have been published in print or on the Web. The few we found most useful are listed in the references. A commonly recommended approach is to start by examining the enterprise's business processes and information assets. In this approach one develops an EIA by first creating common enterprise business and data models [1, 2, 6]. After searching for and failing to find successful examples of applying this approach to comparable organizations we decided that a more productive approach would be to take the approach described in this paper. As a result our contributions are threefold toward developing an EIA for a diverse decentralized organization, developing: (1) a concise representation of its requirements, (2) identifying key elements and a concise iconic representation of their relationships, and (3) an organization structure and process that engages the enterprise as a whole. It is these that are the focus of this paper. More detail can be found in [4].

An EIA represents the coherent framework for putting in place interrelated sets of technology and standards that form the information infrastructure of an enterprise. An infrastructure for digital computing and information has been evolving since LLNL was founded. This *ad hoc* digital infrastructure consists of computing resources, enterprise-wide networking (including Internet connectivity), growing technology for Web use and information repository, e-mail, calendaring and scheduling, computer security practices, desktop computing tools, technology to manipulate large amounts of online data, business and scientific system development matrix organizations, and more. What is missing is an agreed general model of the desired infrastructure and a process to put this in place and assure its coherent and timely evolution.

Like Zachman [6], we were strongly influenced in thinking about what constitutes an EIA by the analogy with building architecture. We are all familiar with the use of the term architecture for the layout and design of a building. A building architecture is the representation of the relationship of all basic components or elements of a building such as rooms, hallways, windows, doors, walls, facades, electrical, heating/air conditioning and plumbing systems, decorative elements, and so forth. The architect uses various types of diagrams, drawings, documents and other artifacts to



schematically represent the structured relationship of all a building's components for its various stakeholders. An architecture must satisfy a set of functional requirements. In the same sense, we define the phrase Enterprise Information Architecture as a structured set of interrelated elements that support all information processes meeting a set of requirements and the representations, documents and other artifacts necessary to make it possible to reach agreement on needed standards and implement and support it over time. Our prime challenge was to develop a concise requirements statement and identify the key EIA elements and their relationships. In the sections that follow, we discuss the elements of an EIA. Our current state and the gap between it and the desired state for each element are discussed in [4].

## **The Enterprise Information Architecture Requirements Represented as Vision, Principle and Strategic Objective Statements**

Any EIA must meet a set of requirements. Our approach was to collect an extensive set of requirements from all the Laboratory's organizations and then boil them down into concise, easy to communicate and understand forms we called vision, principle and strategic objective statements. They are discussed in more detail in [4]. We cannot over emphasize the importance they have played in focusing our attention and getting widespread agreement on the EIA framework developed. Several of these represent significant shifts from current practice. Much work must be done to get from our current state to an EIA that supports these statements.

### ***Vision Statement***

**Easy access to—**

**the right information**

**for the right people**

**at the right time**  
**at the right place**  
**for the right cost.**

### ***Principles Statements***

The seven EIA Principles are:

1. Information is an institutional asset
2. The EIA is the preferred framework for doing business at the Laboratory
3. LLNL-wide access to information is the rule, not the exception
4. The EIA supports ease of use
5. Ownership and stewardship are well defined for all information
6. Information is safeguarded on a risk-defined basis
7. The Enterprise Information Architecture is an evolving framework.

### ***Strategic Objective Statements***

From the Principles, we derived EIA Strategic Objectives:

- All information assets will be integrated when appropriate.
- There will be only one official source for each asset.
- Information capture and validation will be done at its source.
- The information assets shall be readily accessible and available to the ultimate point of use.
- Dissemination, access and user self-service is supported for scholarly, scientific, engineering and business information.
- The safekeeping, storage, retrieval and preservation of our information assets are paramount.

- The extension of information services to the enterprise user's desktop supports the needs of all LLNL staff.
- The extension of information services to all our external customers and partners is desired, where appropriate.
- The continual planning, improvement, and innovation for reliability, availability, serviceability and performance of the information utility is supported institutionally.

## **EIA Standards Philosophy**

Before describing the elements of an EIA, it is useful to say a few words about the philosophy of standards developed. While there was general agreement on the value of standards, there was still concern about having unnecessary standards imposed on organizations. Therefore, the following statement was developed to guide the standard's process.

The Laboratory should set an EIA product, protocol, process, service or other standard, through the EIA stewardship and standardization process outlined in this paper whenever:

- Standardization unequivocally supports/promotes the vision, principle and strategic objective statements
- It supports sharing and interoperability
- A significant/measurable cost savings over current practices can be attained
- An organization is identified as the standard's steward
- An organization is identified and appropriately funded to provide its implementation and operational support (e.g. hardware, software, site licenses, training, system administration, documentation etc.)
- Shared expertise is important.

Standards are not a rulebook, handed down from on high by some autonomous centralized Information Systems authority. Rather, standards are an informed set of decisions, mutually arrived at, that strongly influence information systems practices at LLNL to the mutual benefit of all. With the exception of security, it is expected that occasionally, on a case-by-case basis, different programs and projects may do things in a manner that is in some sense outside of the EIA Standards. These decisions should be based on explicit business considerations that weigh the immediate needs and/or customer requirements of the project against the institutional support and synergy that derive from using the standards.


## **2. Overview of an Enterprise Information Architecture**

### **Pyramid Representation**


As the work of the IAC proceeded, it was critical to identify the key architectural elements and understand their relationship to each other. It was also important to represent them in a way that the LLNL management and staff could easily understand. From this representation we could break up into working groups to develop the current state, desired future state, gap, and recommendations for each element. Our approach worked very well as both a communication tool and representation of important sets of relationships. The EIA is viewed as a structured set of interrelated elements, represented as a pyramid, that support all information processes. The pyramid was chosen because four faces seemed as many as could be easily grasped, it allowed the main elements to be represented, the shape helped emphasize the projected relationships among the faces and layers, and it emphasized the importance of the stewardship activities in the base. As the architecture is fleshed out, detailed requirements are being developed for the various elements. A next level example is shown in fig. 2 for information management. Representing the elements in this way provides a high-level framework for: (1) talking about the EIA and visualizing it; (2) specifying services, standards, and their interactions; and (3) thinking about the organizational structure needed to support the EIA. Until the

IAC reached this representation, it was difficult to get agreement on both what an EIA consisted of and where to focus further work.

The many interacting facets represented by the pyramid faces, base, and interior are all essential.



**Figure 1. Interrelationships of EIA elements: (a) Applications and Layers comprise two faces with EIA Stewardship and Support as the base on which the pyramid rests.**



**Figure 1(b). Top view also shows the Security and System Management faces.**

## **Layers**

Layers represent important elements of LLNL's EIA that build on each other. Each of the application and infrastructure domains represented on the other three pyramid faces requires services from each of the layers: Network, Messaging and Middleware, Information Services, and the User Interface/Desktop. All elements involve data elements, structures and models, which we view as contained in the interior of the pyramid as mentioned later.

### ***Network***

The first or network layer contains three sublayers: (1) all the wiring and fibers and their interconnection topology, (2) electronics for routing and other functions, and (3) the communication protocols necessary to transfer uninterpreted bytes of information between two or more computer systems or devices. This layer includes remote and local communications, including Internet access. Experience suggests that the first two sublayers are particularly important for network reliability, security, ease of maintenance and management, scalability of performance, and cost effectiveness. Network services can only deliver the reliability, security and performance that are provided by the physical infrastructure.

### ***Messaging and Middleware***

In any enterprise network, above the network transport (e.g., TCP/IP), and underneath any specific set of information services or desktop platforms, there is a layer we call the Message and Middleware (M/M) layer. M/M technologies are the glue that provides for transparent, secure information connectivity and integration at several levels of the infrastructure.

The M/M layer has three intimately related sublayers:

1. Information content formats for storage and exchange that are increasingly based on industry standards. Examples include HTML, PDF, RTF, JPEG, MPEG, GIF, MIME, netCDF, the emerging use of the XML/SGML formats, and others. These formats are used in moving interpretable bytes between computers, in displaying (usually interactively) information in

end-user applications, and even as storage formats for many database/data store applications.

The emergence of widely adopted standard content formats contributes significantly to the ability of different applications to interoperate seamlessly.

2. Session, presentation, and application layer protocols are based on industry standards. Examples include SMTP, POP3, and IMAP4 for e-mail, LDAP for directory services, NNTP for threaded newsgroup discussions, HTTP for Web interfaces, SNMP for network management, IIOP for remote service request and replies, X.509 for security using public key infrastructure, JDBC/ODBC for data management, and many more.
3. Common run-time services that support distributed applications, such as directory and naming services, transaction services, security services such as certificate authorities, and the object request brokers (e.g. of the CORBA model) for connecting components. These services depend on many of the protocols listed above, but they are not just protocol specifications; they are implemented software systems running on computers on the network, providing general purpose common services (e.g. naming, security, synchronization and resource management) to various applications, independent of specific clients and servers. In effect they represent a distributed operating system.

In broad terms the network layer transmits uninterrupted bytes between end points, while the Message part of the M/M layer encompasses (upper layer) protocols and formats that convey meaning to these bytes, and the Middleware part refers to common run-time services that exist as part of the enterprise architecture, and are useful to many applications.

### ***Information Services***

It is assumed in the EIA that both a standalone set of desktop services and a client/server architecture are supported. A server role provides services; a client role makes service requests of the servers. The Information Services layer of the architecture consists of the server side of applications. Examples of services at this layer of the architecture are the POP e-mail servers, HTTP-based Web servers, and



Meeting Maker calendaring and scheduling servers. Other examples are database and document servers, and servers that support business operations and programmatic applications. The greater the extent to which we can standardize on services at this layer that are built on industry standard APIs and protocols, the wider will be our choices of vendors and interoperability both inside and outside the Laboratory.

Further, to the extent that business operations and programmatic applications are built on EIA and industry standards, the quicker these services will be made available and the lower the cost. The information services layer and associated applications are driven by both enterprise-wide and specific business and programmatic requirement. While we recognized the importance of software engineering methodologies and development environments to creating information services and applications they were viewed as outside the work of the IAC. Two important Application Domains with broad Labwide use, Collaboration and Information Management—and the Information Services layer associated with them—are discussed below. Even though details of what is in this layer are driven by the other pyramid faces, its existence and the need for corresponding standards needs to be made explicit in the model.

### ***Desktop (User Interface)***

The Desktop layer represents the user interface with the rest of the architecture. We broke the user interface discussion into general desktop support and Web issues (in general the Web technology spans all the layers). This layer includes whatever system is most appropriate to access needed information services (e.g., laptops, mobile hand-held devices, or possibly a kiosk for users who do not normally work with computers). While the services provided at the lower layers of the pyramid are important to users, they are generally maintained by others. They can be viewed like a building's heating/air conditioning, electricity, or plumbing infrastructure. Occupants want them to work but prefer minimal knowledge and involvement. The Desktop layer contains the hardware (e.g., PCs, workstations, laptops, or X-terminals) that gives direct contact between the user and the services, the desktop system software, standalone applications such as word processors or spreadsheets, and client

applications used to access or manage server-based applications such as data access, e-mail or calendaring. A Web browser is an important client because of its increasing use as the standard information service user interface.

Web technology is a highly efficient mechanism with which to publish and access information internally and worldwide. The Web-enabled desktop (Webtop) offers the potential for a standardized user interfaces for both scientific and business applications, whether the application is purchased or developed in-house.

## **Application Domains**

The applications side of the pyramid represents all applications of interest. Important classes of intersecting applications include collaboration tools such as e-mail and calendaring and scheduling, information sharing tools such as document repositories or databases, and all the business and programmatic (scientific and engineering) applications. To understand how all the layers work together, consider two concrete examples: e-mail and online document access. Some of the various clients, servers, and protocols of interest at each layer are shown in Table 1.

Any given application requires support from, and places requirements on, the protocols and services in all the layers. The greater the extent to which industry standards are adopted in the layers, the greater flexibility there will be to choose products from multiple vendors. For example, standard Internet protocols in the Network and Messaging layers allow us to obtain compliant post office servers and e-mail clients from multiple vendors for Laboratory e-mail service. However, to minimize support costs, limits may be placed on the number of products that will be supported as an LLNL standard.

**Table 1. Standards in EIA layers for collaboration and information access applications.**

Layers	Applications	
	E-mail ( <i>Collaboration</i> )	Online Access ( <i>Information sharing</i> )
Desktop	Eudora Clients	Netscape Browsers
Information Services	Pop Servers, Ph servers	Online document repositories in TID
Messaging	SMTP, POP3, IMAP, MIME	HTTP, HTML, PDF
Network	TCP/IP, Internet and LabNet	

We recognize that not all the requirements for business or programmatic applications can be met by standard services. However, as the LLNL EIA grows and matures, we believe that an increasing base of services will be available to make it easier, less costly, and less time consuming to develop such nonstandard applications. Two application domains were identified with particular importance to the EIA vision collaboration and information management. These are discussed in more detail in [4].

### ***Collaboration Applications Domain***

Computer-aided collaboration support, also called groupware, is the broad, somewhat fuzzy area consisting of computer systems that enable two or more individuals to work together more effectively.

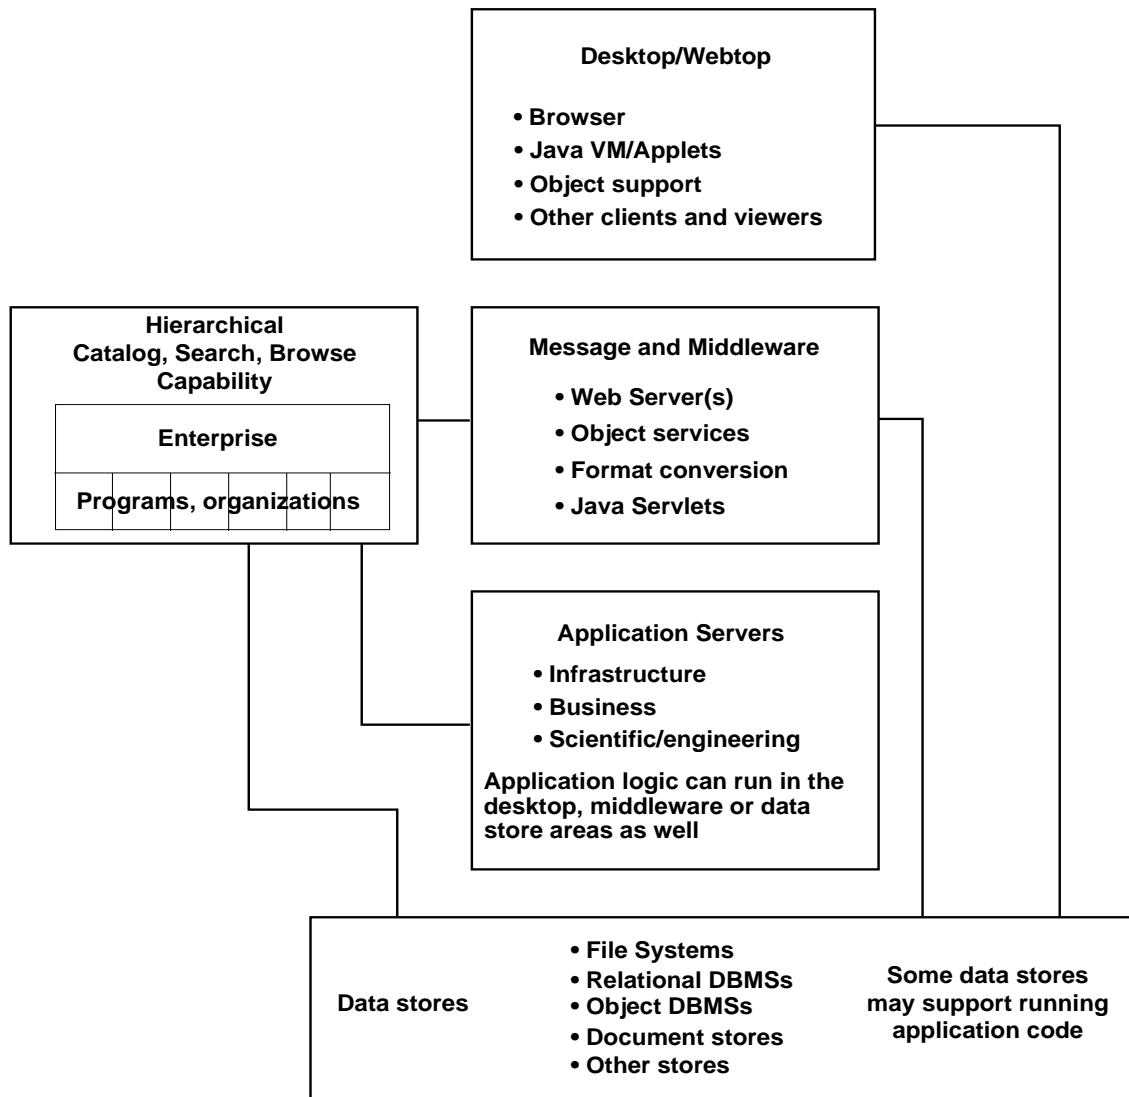
Electronic mail is the most widely understood kind of groupware, but there are other groupware paradigms identified for future study shown in Table 2:

**Table 2. Groupware paradigms.**

	Same Place	Different Place
Same Time (synchronous)	Meeting Support Systems	Chat Shared Whiteboard Audio/Video Conferencing
Different Time (asynchronous)		Electronic Mail (e-mail) Calendaring and Scheduling Electronic Forums (discussion lists, newsgroups) Multiple Authorship Support Document Management, Electronic Library Workflow applications (e.g., purchasing)

### ***Information Management Application Domain***

The information management application domain is an important driver for an Enterprise Information Architecture. Scientists are estimated to spend 20-40% of their time searching for and gathering information and we expect the percentage is similar for those who work with business and other information. Clearly other applications need to leverage or connect to the capabilities for information management. We examined several information management systems both within and without the Laboratory and abstracted the high level model shown in Fig. 2 to focus standards activities.



**Figure 2. High-level Information Management Model.**

With this model we identified areas in the desktop, information services ( e.g. applications and datastores), M/M, and security areas where enterprise standards would be helpful. Of particular importance is the need for an enterprise wide catalog system that allows other catalogs developed independently to be integrated. We also identified developing an enterprise inventory of information assets as a high priority task. Details of this model and LLNL information management systems that point the way toward such an integrated system are described in [4].

## Unclassified Security

At LLNL all classified information and associated processing and other resources are on a separate network, which is not connected to the outside world. The architecture considered here is for unclassified information.

The Security/Privacy side of the pyramid spans all the layers and interacts with the Applications and System Management sides. The architecture needs to provide perimeter defenses, defenses in depth on local area networks and end systems, as well as provide detection and reaction systems to handle possible external attacks and breaches of perimeter defenses. To develop a detailed security/privacy architecture, we need to understand the information and other resource assets that need to be protected, the threats to these assets, and the technologies available to counter these threats. Consider two examples requiring security-related standards snooping and authentication/authorization.

Snooping on the network is possible if a bus-style connection topology (classical Ethernet) is used because all information is visible to all computers on the network. A star point-to-point wiring topology connected at the hub by a switch removes this visibility. This topology and switching must be specified at the Network layer and can be supported, for example, by switched Ethernet or ATM technologies.

Authentication and authorization are other important security services as part of the above and to protect proprietary information of many kinds. They must be supported with a secure network topology, appropriate servers for generation and distribution of certificates, and standard message formats for certificate exchange and validation. User applications and client software on the desktop, such as the Web browser or e-mail client, also must be enabled to generate, accept and check these certificates. Other services at the desktop, such as secure networking capabilities, may be required for safe input of passwords. Thus, a secure authentication and authorization system requires services at all layers of the architecture, including applications at the information service and desktop client levels.

## **System Management**

The final face of the pyramid shown in Fig. 1(b), System Management, similarly requires services from the layers and interacts with the Applications and Security faces. It represents all the services needed to manage, troubleshoot, and debug applications as well as to maintain the services specific to a given layer. For example, consider automatic distribution services for keeping the software on the desktop up to date. This requires client mechanisms in the user interface, software repository and updating services, agreement on message formats, security mechanisms, and so forth. Within a given layer, such as the network layer, system management provides the facilities needed to detect and isolate the wiring and electronics faults, keep routing tables up to date, support the movement of staff to new office locations, and many others.

## **Human Activities—the Base of the Pyramid**

The base of the pyramid, shown extended, represents the human activities required to support all aspects of the EIA, such as stewardship of its evolution, day-to-day operation of all its services, and its use. Setting up an organization and processes for enterprise wide interaction and stewardship issues is crucial for long term success and discuss this aspect in more detail below.

## **Data and Information—the Interior of the Pyramid**

Data and information are required in order to support all aspects of effective and efficient use of information—generation, storage, access, presentation, sharing, manipulation and transformation, interchange, management/administration, and security. In terms of Fig. 1, we can visualize information as residing within the pyramid. System data and metadata is associated with every layer, while data and information the user directly cares about (such as e-mail, documents, or business/programmatic information) are managed by the information services layer, created and interpreted by applications and accessed by the user interface/desktop layer. It is within this volume, associated with the appropriate layers and sides that various data architectures need to be developed. These will consist of standard data types, structures and models useful across the enterprise.

### **3. Enterprise Information Architecture Stewardship Organization and Processes**

The EIA Stewardship process and organization shown in Fig. 3 are discussed below.

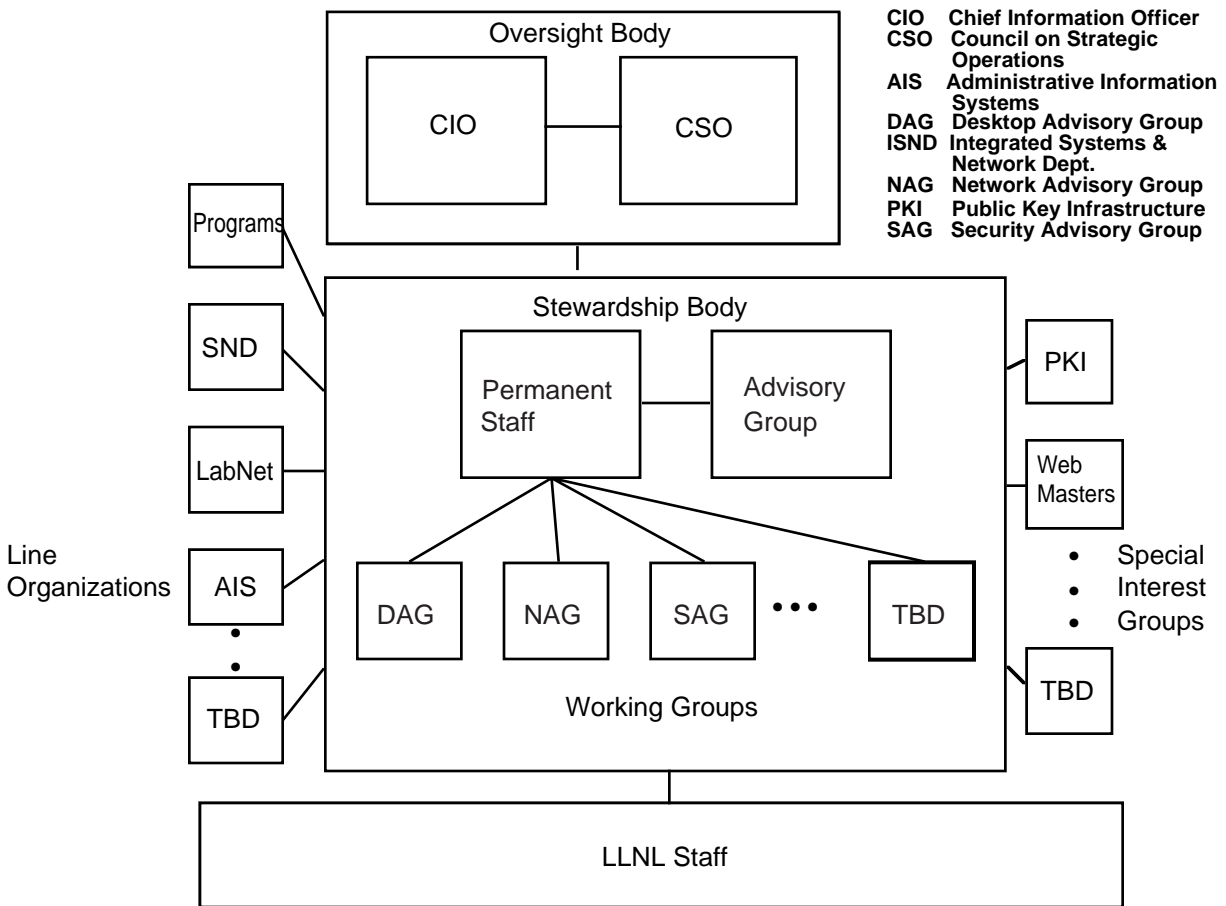
#### **EIA Stewardship Organization**

Because the EIA is a Laboratory-wide activity, there are a number of roles and organizations involved in its stewardship shown in Fig. 3. The organization names in the Line Organization, Working Group, and Special Interest Group boxes are representative of the types of groups involved. There are two main roles in standards stewardship: (1) stewardship of the standards themselves, which is the responsibility of the Stewardship Body, and (2) implementation and operational support, which is the responsibility of the line organizations.

#### ***EIA Stewardship Body***

The Stewardship Body consists of a small permanent staff (~4 people), an EIA Advisory Group representative of EIA stakeholders from across the Laboratory, and a set of permanent and temporary working and technical advisory groups responsible for EIA standards stewardship and special studies. The charter of the Stewardship Body is to provide the stewardship of EIA standards and process and manage the EIA activity on an ongoing basis.





**Figure 3. EIA Stewardship Organization.**

#### The Stewardship Body:

- Assumes stewardship of the EIA Vision and Principles. Should these need updating in the future, the Stewardship Body would coordinate that process with the Oversight Body and LLNL staff, and might utilize the RFC process before adopting modifications.
- Coordinates and drives the process of proposing, refining, and adopting the architecture and its associated body of standards as they evolve.
- Interacts with LLNL staff and management on a regular, ongoing basis to ensure the EIA is meeting business and programmatic needs; fosters the exchange of ideas and information in both directions; and coordinate tech watch activities, such as gathering information on emerging technologies or conducting limited-scale pilot/proof-of-concept projects.

- Interacts with LLNL's institutional information service providers who provide the day-to-day operational support of the EIA and related activities to ensure the exchange of ideas and information and to promote Labwide cooperation and leveraging of one another's activities.

When a standard or other EIA service is established, a working group needs to be identified or formed for its ongoing stewardship.

### ***Oversight Body***

The EIA Oversight Body depicted in Fig. 3 is composed of the CIO, and other senior Laboratory managers on the Council for Strategic Operations. The Laboratory CIO has the senior management responsibility for the EIA. This includes activities such as final concurrence or approval for all EIA standards, strategic direction, and budget. The Council for Strategic Operations provides an advise and consent role to the CIO on high level EIA policy, direction, and funding.

### ***LLNL Staff***

The LLNL staff is made up of all Laboratory employees. All members of the staff can participate in the Stewardship process by exchanging ideas with the Stewardship Body about any aspect of the architecture. Although a great deal of technical expertise resides in the Stewardship Body, particularly the working groups, we recognize that much real work in information systems development and utilization is performed in the Laboratory's organizations and projects. Customer perspective, customer needs, and appropriate input regarding cost effectiveness can only come from the staff at large.

### ***Line Organizations***

The implementation of the standards and other EIA services and their day-to-day support is the responsibility of line organizations such as those shown in Fig. 3. When a standard or other EIA service is established, one or more groups need to be identified and funded for its implementation and ongoing support.

## ***Special Interest Groups***

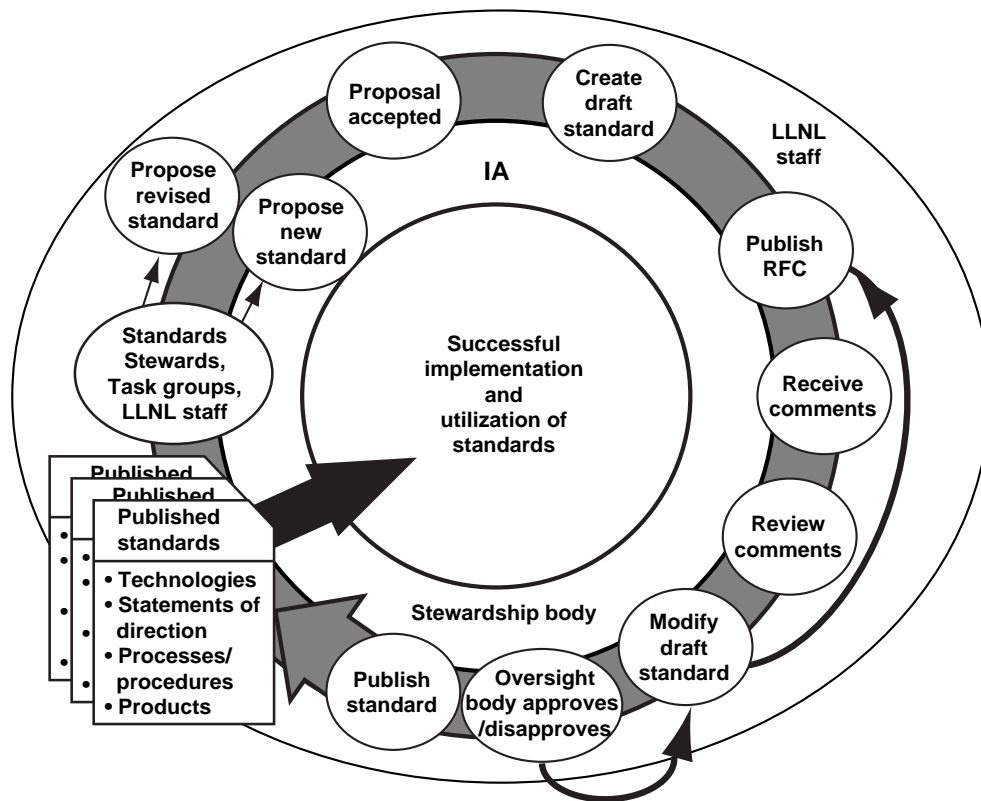
Special Interest Groups are constantly forming and dissolving to deal with issues around new areas of technology and share experiences. This process is encouraged. Because these groups are sources of expertise and suggestions for additions and modifications to the EIA and standards, the Stewardship Body works with such groups.

## **The EIA Standards**

EIA Standards are of three main forms: (1) Statements of Direction provide institutional guidance toward a common goal; they are useful in procurements and system design and integration decisions, when it may still be too early to put technology or product standards in place; (2) technology and product standards that may include both industry and product standards; and (3) procedure or process standards. It is assumed that policies for security or other functions may also employ the processes Fig. 4 if appropriate.

## **Standards Process**

The EIA Standards Process depicted in Fig. 4 is a crucial element of the Enterprise Information Architecture. It is the means by which informed decisions are made and broad consensus reached on standards that drive Laboratory information practices. Standards emerge in two ways. First, there could be an exploratory phase in which ideas are exchanged, and experimental/prototype/technology demonstration activities are



**Figure 4. The EIA Standards development and review process.**

conducted. Alternatively, new or updated standards may be proposed through other mechanisms—the steward of a standard may put forth a revised version, or any staff member can propose a standard at any time. In any event, when a new or revised standard has been proposed, the Stewardship Body will examine it. The Stewardship Body can reject the proposal because it fails to meet one or more of the criteria for standards listed earlier.

If a proposed standard is judged worthy of further consideration it is passed through a Request for comments (RFC) process, used successfully for almost three decades in the Internet community. A formal RFC activity is conducted for each proposed standard (Fig. 4). Comments from the staff are collected and reviewed to determine whether a proposed standard can go forward for adoption, requires significant revision and more review, or should be discarded. The proposal is advertised through several Labwide channels and is made available on the EIA Web pages. The RFC Review

Team will issue a roll-up of the review, summarizing the input received. The entire Stewardship Body then determines the next step, which will be one of the following:

1. Reject the proposal altogether
2. Return the proposal to its source for revision
3. Recommend the proposal to the Oversight Body as an official Enterprise Information Architecture Standard (or recommend adopting it with minor modifications).

The Oversight Body will accept or reject the recommendation of the Stewardship Body, possibly leading to an iteration of one or more of the steps above.

## **4. Summary**

We have described a high level framework for an EIA and the support organization and process that we think are appropriate for many organizations, particularly for those with decentralized, autonomous cultures and organization structures. The three key elements include: (1) capturing the requirements in concise, easy to understand vision and principle statements, (2) identifying key EIA elements and their interrelationships and in an iconic diagram, and (3) specifying an enterprise-wide organization structure and processes for the EIA's development, implementation, support and ongoing evolution.

Our status is that the EIA has been widely discussed with staff and management (although much more is required); funding has been obtained to launch the Stewardship Organization; active task and advisory groups have been established for the Desktop, Web, Network, Security elements; initial explorations have been started in the information management area for catalog metadata and document management standards; and Desktop, collaboration, Web and related M/M and Information Service standards have been put in place. Implementing the unclassified security architecture recommended by the Security Task Group has very high priority. Much work remains to be accomplished to realize the EIA described in this paper and [4]. We recognize that the funding, development, implementation and cultural change required will take several years.

## Acknowledgements

The work reported here is the result of the contributions of over 60 people from across the Laboratory who participated in the IAC and the task groups associated with each element shown in Fig. 1. These people interacted with many more within their organizations. The chairs of the task groups deserve special mention for pulling together the current state, desired future state, gap and recommendations for each element: Dave Wiltzius – Networks and Stewardship processes, Rob Allen – Collaboration and Message/Middleware, Barbara Campbell – Desktop, Ted Michels – Web, Barry Howard – security, Dave Seibel Vision and Principles. Carol Hunter played an important role as Cochair of the IAC until other priorities claimed her attention. We want to acknowledge our colleagues at Los Alamos National Laboratory (LANL) whose EIA work was of most influence on ours [5], particularly the helpful interactions with Diane Weir the LANL EIA project manager. The work reported here was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract W-7405-ENG-48.

## References

1. Cook, Melissa, Building Enterprise Information Architectures, Prentice Hall 1996.
2. DOE, *Enterprise Information Architecture Volume I. The Foundations*, March 1995, <http://www-it.hr.doe.gov/IAt/Iapro>
3. Jet Propulsion Laboratory, *Information System Architecture*, JPL D-12991, Nov. 1, 1995.
4. LLNL Information Architecture Report UCRL-AR-130900, July 25, 1998  
<http://www.llnl.gov/projects/IA/>
5. Los Alamos National Laboratory, *Information Architecture: The Foundation*, LA-UR-94-1433, May 3, 1994. <http://www.lanl.gov/projects/IA/welcome.html>
6. Zachman, John A., “A Framework for Information systems Architecture,” IBM Systems Journal, Vol. 26, No. 3, 1987.

